Tuning Crowdsourced Human Computation

Chen Cao[§] Zheng Liu [§] Lei Chen[§] H. V. Jagadish [†]
[§]Hong Kong University of Science and Technology, Hong Kong, China
[†]University of Michigan, Ann Arbor, MI, USA
caochen@cse.ust.hk, zliu@cse.ust.hk, leichen@cse.ust.hk, jag@umich.edud

arXiv:1610.04429v1 [cs.HC] 14 Oct 2016

ABSTRACT

As the use of crowdsourcing increases, it is important to think about performance optimization. For this purpose, it is possible to think about each worker as a **HPU**(*Human Processing Unit* [1]), and to draw inspiration from performance optimization on traditional computers or cloud nodes with CPUs. However, as we characterize HPUs in detail for this purpose, we find that there are important differences between CPUs and HPUs, leading to the need for completely new optimization algorithms.

In this paper, we study the specific optimization problem of obtaining results fastest for a crowd sourced job with a fixed total budget. In crowdsourcing, jobs are usually broken down into sets of small tasks, which are assigned to workers one at a time. We consider three scenarios of increasing complexity: *Identical Round Homogeneous* tasks, *Multiplex Round Homogeneous* tasks, and *Multiple Round Heterogeneous* tasks. For each scenario, we analyze the stochastic behavior of the HPU clock-rate as a function of the remuneration offered. After that, we develop an optimum Budget Allocation strategy to minimize the latency for job completion. We validate our results through extensive simulations and experiments on Amazon Mechanical Turk.

Keywords

Crowdsourcing; Algorithm

1. INTRODUCTION

Human Computation [2] has emerged in recent years as a new and exciting compute paradigm. As a powerful compement of traditional computer systems, human computation naturally allows tasks with human-intrinsic values or features, like comparing emotions of speeches, identifying objects in images and so on. The emergence of public crowdsourcing platforms, which provide a scalable manageable workforce resource, has boosted the utilization of this long-discovered [3] human cognitive ability. A wide range of data-driven applications now benefit from human computation by considering it as a new computing component. Examples include *a*) crowd-powered databases [4–6] and fundamental operators like filtering [7] and Max [8, 9], group-by [10], *b*) advanced data processing technologies like image tagging [11], schema matching [12] and entity resolution [13], and *c*) combinatorial problems like planning [14] and mining [15].

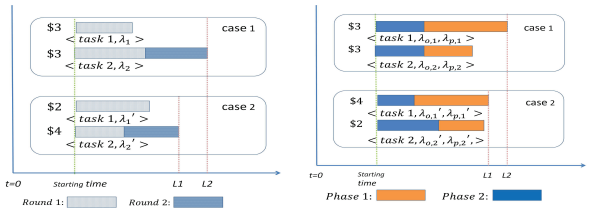

(a) Example 1: Repetition (b) Example 2: Heterogeneous

Figure 1: Demonstration of Motivation Example

As crowdsourcing becomes more prevalent, there is an effort to understand and characterize it better. In this regard, it has been suggested that the system can be viewed as comprising Human Processing Units (HPUs) that are analogous to CPUs of traditional computers. As atomic task performed by a worker is then one "instruction" of the HPU, and the time to respond is the "clock cycle". However, the HPU has many characteristics that differ from those of a CPU:

i) the clock time is **stochastic**; *ii*) the results are **error-prone** according to a probability; *iii*) the cost includes **monetary** expense.

Given the HPU abstractions, one can consider optimizing many aspects of HPU processing. Our focus, in this paper, is the HPU clock rate. This is because we want to minimize the total latency of a computational task by optimizing HPU clock rate.

The clock rate for the HPU is variable, and is different for each instruction and each instance. In a crowdsourcing workforce market, a task is exposed on the market with a promised reward, and then "workers" select the task to work on according to their interests. Recent studies on Amazon Mechanical Turk (AMT) report that the task acceptance duration follows an exponential distribution [16, 17], and the rate λ_i is mostly determined by the promised reward [18] and the type(difficulty) of the task [17]. Then, the processing time of a task follows an exponential distribution with another rate [19], which is independent of the promised payment [18].

Once the task has been specified, the only task-owner input that can control the completion time is the *payment*. If we only have one task to be performed by HPUs, the solution is very simple – the more we can afford to pay, the faster the task will be completed.

Of course, the computational job at hand is typically performed with the aid of many HPU tasks [4–6]. In fact, a typical algorithm architecture repeats each task in multiple times. Thus, the requester issues a large number of HPU tasks in parallel, each possibly to be repeated, and then waits for all HPUs to return result. As such,

ACM ISBN 978-1-4503-2138-9.
DOI: 10.1145/1235

there is limited value to optimize the clock rate of a single HPU in isolation: what really matters is the latency of the entire computation, which is determined by the longest duration among the set of parallel repeated tasks. Thus, our optimizing HPU clock rate problem is in fact on studying how to allocate a given fixed budget B in a manner that minimizes the total latency of a computational job involving HPUs.

To demonstrate the challenging issues involved in optimizing HPU clock rates, let us consider the following two motivating examples, both based on a crowd-powered data-base, as proposed in [4–6, 8, 9].

MOTIVATION EXAMPLE 1. *(Figure 1(a)) Consider a sorting task on 4 given items $O = \{o_1, o_2, o_3, o_4\}$. According to the user's requirements, the query planner, for example the "next votes" proposed in [9], decomposes the sorting task into atomic pairwise voting tasks $T = \{\{o_1, o_2\} \times 1, \{o_3, o_4\} \times 2\}$, which means the HPU is expected to run the task of comparison on such pairs for 1 and 2 repetitions (times) respectively. As illustrated in Figure 1(a), the two tasks commence at the same time, but in order to finalize the entire query, the database has to wait until the end of the longest atomic task. There are many choices for budget allocation. Two obvious ones are: one is evenly divided to two tasks, 3 for task 1 and 3 for task 2(case 1); another one is more load-sensitive, 2 for task 1 and 4 for task 2(case 2). The results for the two cases are shown in the figure, suggesting that the second option is better. But how could we predict this? Moreover, even if this is the better of these two choices, is it the best? What is the best allocation across the two tasks?*

MOTIVATION EXAMPLE 2. *(Figure 1(b)) Consider now a more complex scenario in which the database is required to process two types of queries simultaneously, sorting and filtering [7], where the latter can also be decomposed into pairwise voting tasks(yes or no voting). Suppose two tasks are given $T = \{\{o_1, o_2\} \times 1, \{o_3, yes?no\} \times 1\}$. However, unlike the previous case, different types of task present different difficulty levels, which leads to different rates. As shown in Table 3, for the same price p_i, the processing rate λ_p of sort voting is lower than that yes or no voting. In addition, the entire latency depends on both how long a task is offered before it is accepted, which depends on the reward offered, and also how long it takes to complete the task, which depends on the task itself but not on price. Unlike in the previous example, we now have to take such latency into account as well. Trying two budget allocations: evenly allocating \$3 to two tasks; balancing budget according to difficulty, sorting task with \$4 and filtering task with \$2, we get latencies as shown in the figure. Once again, these are obviously not the only allocations possible, and our interest is in finding the optimum, compounded by the difficulty of predicting the uptake rate for any reward level (notice that Table 3 only gives us values for a few price points), and of additionally folding in the task completion time into the framework.*

As shown in the motivating example above, the allocation of budget to the tasks matters to a great extent in terms of the overall HPU processing latency. The difficulty of finding the optimal allocation strategy is two-fold: 1) the latency of an atomic task is a random variable which depends on the type of the task, the allocated budget, and the current workforce market situation, therefore it is non-trivial to predict the overall latency of a set of tasks, particularly when they are of different types; 2) the search space for finding an optimal solution is large so that efficient algorithms and/or approximation tradeoffs are necessary. The promised payment has a minimum granularity(\$0.01 on AMT), which renders the tuning process a discrete, rather than a continuous optimization problem.

To address these challenges, the following contributions are made:

- In Section 3, we begin with HPU characteristics, develop a stochastic model to predict uptake rate as a function of reward amount, and show how to estimate the model parameters. Using these results, we can determine the expected latency for any specific budget allocation choice.
- In Section 4,we formally propose the *H-Tuning Problem* to minimize the expected latency of a given set of tasks and propose probabilistic analysis and tuning strategies under three practical scenarios: *Homogeneous, Repetition and Heterogeneous*. In each case, we show how to solve an optimization problem with a large feasible space of possible budget allocation choices.
- The performance of proposed strategies are verified on real crowdsourcing platform, and with simulation in Section 5.

In addition, Section 2 gives an overview of the related work. Section 6 makes conclusion to this work.

2. RELATED WORK

Leveraging the HPU in hope of better performance is an attractive topic ever since the emergence of crowdsourcing applications. Many recent works have studied various optimization issues associated with the HPU [20,21]. Most of them focus on the quality issue in terms of answer confidence [22,23], and some efforts on the optimization of monetary cost [17].However, in the effort of designing an industrial level computing module, speed or latency is always one of the most significant concerns among the various properties. Most of current works touch this issue by reducing the number of queries issued to the crowds [2, 6, 7, 9, 12, 14, 15].Whereas, considering the HPU a new "hardware" for general human computation, a lower-level clock-rate model, instead of the higher-level number of queries, is far more entailed. Unfortunately, the stochastic human behavior makes this model rather intractable. Several applications tried to optimize the HPU's performance in real time in order to finish tasks before a preset deadline [19, 24, 25]. But their approaches are highly application-dependent and thus hard to adapt to a general framework; in addition, the "deadline" semantic does not support the batch processing scenario where a general HPU usually meets.

Meanwhile, another practical methodology is developed by recruiting a set of prepaid worker, so that they can wait online and process the task immediately after publishing. In the work of [26], the authors propose such a pre-paid model to instantiate a real-time respond crowdsourcing interface, and a Retailer Model is adopted to describe the prepaid workers behavior [27]. Following the Retainer Model, one analytic effort on optimally organizing the microtasks can be found in [28]. Note that the prepaid implementation differs greatly from this work: the tasks for prepaid implementation entails high instantaneity, where the tasks are expected to be finalized in several seconds(the payments are relatively higher as well); however, the HPU tuning assumes a system-level perspective, where the latency of the task set varies over a larger range according to the specific requirement of the database users. Last, the *Queuing Theory* based model of prepaid implementation cannot be tailored into the HPU scenarios easily.

This work is most related to [29], where the problem of minimizing crowdsourcing latency is formulated into two optimization issues: *1. minimizing the completion cost of all the tasks given deterministic deadline of every task*, and *2. minimizing the latency with constrained budget*. The objective of the problem discussed in this work is virtually same with the second issue in above work.

Table 1: HPU Processing Rate for Motivation Example

reward($) \ task type	sorting vote	*yes* or *no* vote
2	2	3
3	3	5
1.5	1.5	2

However, our work is distinguished with [29] in the following aspects. First, the latency of a crowdsourcing task is modeled with two phases: the on-hold phase and the processing phase. Such consideration is consistent with the real world scenarios. However, [29] only considers the latency of the tasks' acceptance. Secondly, the crowdsourcing tasks can be processed both parallel (multiple tasks being processed simultaneously) and sequentially (one task calls for multiple answering repetitions, which are submitted one after another). Both processing manners are studied in this work, while [29] minimizes the latency with the implicit setting of pure parallel processing.

3. THE HPU MODEL

In this section, we begin with the basic crowdsourcing framework, develop the HPU model and demonstrate how to estimate HPUs' parameters. In short, we lay the foundation for the optimization problem we consider in the next section.

We begin with definitions of standard crowdsourcing concepts:

- **Requester**: A requester publishes tasks, collects answers, and makes the promised reward payments. Database-wise, the *requester* is the higher-level *"executor"* as in [6] or *"task manager"* in [5]. The requester has also been called the"task-holder", "job-owner" or "builder".

- **Worker**: A worker (or crowd-worker) performs the actual human processing tasks. A worker arrives at the market in a uniformly random manner, and she immediately chooses one of the tasks to work on. The preference of task selection is based on her utility maximization principle. After a period of time, the worker finalizes the task by returning the answer to the requester. Note that some research [17] reports that the worker activity on Amazon MTurk observes fluctuation along both a daily and a weekly basis. However due to the scale of data-driven micro-tasks, which are mainly light-weight voting, such long-term fluctuation can be ignored, provided that we use parameters that recurrent. In Section 3.3 we discuss the practical methodology to infer the realtime system parameters.

- **Task**: A task is the most decomposed operation that a **worker** may work on. Unfortunately, there are intrinsic limits of human cognitive capacity [2], and huge differences are observed in the demographics of crowd workers [30]. Consequently, to ensure the coherence and reliability of the human answers, a **worker** is restricted to perform a set of most basic operations like selecting from several options, ranking within a couple of objects, connecting between figures, tagging images with text and so on. Many of these human operations can be categorized into voting, where a latent true option needs to be located with some effort (a period of time).

  In the literature, tasks have sometimes been called "jobs", "HIT(Human Intelligent Tasks)" and so on. However, we reserve the word "job" for the following:

- **Job**: A job is what the requester is responsible for. A job is accomplished by invoking tasks in parallel in one or more phases, with possible additional computation performed at the requester at the beginning and end of each phase. In this paper, we will consider three different structures for tasks in phases, as we shall see below.

3.1 Worker Selection Model

Based on the definition of *worker* above, in a workforce market, a worker appears and starts working on a task uniformly at any time. Meanwhile, the worker's preference among the candidate tasks relies on the subjective utility measurement.

3.1.1 Worker Appearing Time

The online workers enter the crowdsourcing market with a random manner. For a short period of time, like a few hours for platforms like Amazon MTurk, (according to the statistics of workers' arrival which is publicly released on AMT), the workers' arrival rate (the number of workers arrive within the unit time) can be regarded to be a constant number. Such property enables us to model the worker's appearing time with the following process. Denote the current workers' arrival rate with the constant number λ. For a time interval of fixed length Δt, the probability of *No worker appears* equals to $(1 - \Delta t \cdot \lambda)$. Suppose a task is submitted at time "0", and the task is accepted rightly after a worker arrives, the distribution of its acceptance can be derived as follows: $P(t_{acc} \leq s) = 1 - P(N(s) = 0) = 1 - (1 - \Delta t \cdot \lambda)^{\frac{s}{\Delta t}}$,, where t_{acc} is the time when the task is accepted and $N(s)$ denotes the number of arriving worker at time stamp s. Taking limit to Δt gives the following expression: $P(t_{acc} \leq s) = 1 - \lim_{\Delta t \to 0}(1 - \Delta t \cdot \lambda)^{\frac{s}{\Delta t}} = 1 - e^{-\lambda s}$. Clearly, the acceptance time of a task follows exponential distribution on condition that the task is accepted once a worker arrives.

Recent research work in [16, 17] delve into more detailed analysis of when workers appear, with more delicate consideration of time period and so on. Nevertheless, for an encapsulated computation module, a major exponential model is powerful enough for describing the latency characteristics.

3.1.2 Task Preference

In previous discussion, we make the assumption that a task is accepted once a worker arrives. However, workers have preferences over the tasks and tend to choose the task that can maximize her benefits. In other words, a task is accepted by an appearing worker with certain probability "p". Since we have pointed that after the submission of a task, the latency can only be adjusted though pricing, therefore p is set to be variable affected by the task's price "c" ("$p(c)$"). Together with the worker's arrival rate, the probability of *No worker accepts a task* is derived as: $(1 - \lambda p(c)\Delta t)$ (when $p(c) = 1$, such expression is equivalent to probability of "No worker arrives" presented in last part). Following the same procedure, the task's acceptance distribution is re-formulated as: $P(T < t) = 1 - e^{-\lambda_c t} = 1 - e^{-\lambda p(c) t}$, where λ_c ($\lambda_c = \lambda p$(c)) is the joint acceptance rate of price c.

A detailed discussion of choice model can be found in [17]. But to better estimate the latency behavior, in Section 3.3, we present a real-time technique to infer parameters for tuning strategies.

3.2 The HPU Latency

Like in traditional CPU-based applications, when a single task is published to the HPU, there will be two phases before the answers are returned and collected: on-hold phase and processing phase. The first one is the period from the task being published to the task being chosen by a worker; the second one is the period waiting for

Table 2: Summary of Notations

Notation	Description
t_i	an atomic task
T_N	a set of atomic tasks with size N
$L(t_i)$	latency of task t_i
K	the maximum number of possible batches
$L_o^{b_i}$	*On-hold* latency of batch b_i
$L_p^{b_i}$	*Processing* latency of batch b_i
λ_o^i	the *On-hold clock rate* of batch b_i
λ_p^i	the *Processing clock rate* of batch b_i
B	the total budget
g_i	task group i, whose tasks are of i repetitions
p_i	payment for the task group i
$E(g_i)$	the expected latency for the task group i

answer from the worker. Statistical research has been conducted on several crowdsourcing platforms to capture the traits of such latencies [16–18].

DEFINITION 1 (LATENCY). *The On-hold Latency L_o of a task (or a batch of tasks) is the clock time from when the task is published to the time when it is accepted by a worker. The Processing Latency L_p of a task (or a batch of tasks) is the clock time from when the task is accepted to the time when the answer is returned and collected by the system. The Overall Latency L is the sum of L_o and L_p: $L = L_o + L_p$.*

According to the worker appearing behavior proposed previously, we can derive that the distribution of the overall latency as follows. Let λ_o and λ_p denote the *clock rates* of the process in On-hold and Processing phase respectively, and the probability density function of the latencies are as follows

$$f_o(t) = pdf(L_o \le t) = \lambda_o e^{-\lambda_o t}, f_p(t) = pdf(L_p \le t) = \lambda_p e^{-\lambda_p t}$$

Since the latency of On-hold phase depends on the attractiveness of a task towards the crowds, whereas the latency of Processing phase depends on the actual cognitive load of a task, we assume these two phases are independent from each other, which is supported by a recent study [19]. Therefore, the probability density function for the overall latency L can be derive as following.

$$f_L(t) = pdf(L \le t) = f_o(t) * f_p(t) = \int_0^t \lambda_o e^{-\lambda_o(t-u)} \lambda_p e^{-\lambda_p u}\, du$$
$$= \frac{\lambda_o \lambda_p}{\lambda_o - \lambda_p}(e^{-\lambda_p t} - e^{-\lambda_o t}) = \frac{\lambda_o \lambda_p}{\lambda_o - \lambda_p}(e^{-\lambda_p t} - e^{-\lambda_o t}),$$

where "$*$" denotes the convolution operation of two pdf.

3.2.1 Parallel Processing

In order to complete tasks quickly, unrelated tasks will be published simultaneously onto the crowdsourcing platforms. Given a set of k batch tasks, B_k, being processed by the HPU simultaneously, the distribution for the overall latency of parallel processing is the maximum latency of all the tasks:

$$F_{para}(t) = cdf(L_{para}(B_k) \le t) = \prod_{i=1}^{k} cdf(L(b_i) \le t)$$

EXAMPLE 1. *Revisiting the motivating examples of the introduction, we now have the machinery in place to discuss how we obtained the latencies shown in (Figure 1(a) and (Figure 1(b). The expectation of the longest task for the first example is*

$$E[L] = \frac{1}{\lambda_1 + \lambda_2}(1 + 2\frac{\lambda_1}{\lambda_2} + \lambda_2(1 + \frac{\lambda_1}{\lambda_2} + \frac{\lambda_2}{\lambda_1}))$$

Based on Table 3, $E[case_1] = 2.93(s)$ and $E[case_2] = 2.25(s)$, where the load-sensitive strategy is better.

A similar computation for the second example shows that the expected latency becomes 3.5s and 2.7s respectively.

3.3 The HPU Running Parameters

The crowds workforce platform is always fluctuating, both in terms of demographics and in population. However, an exponential model suffices as a good approximation. To support a robust tuning strategy, we propose to statistically infer the parameters with following two methodologies.

3.3.1 Parameter Inference

To infer the parameter λ_o, a "probe" program is introduced, which publishes tasks with varying prices. The workers who accept the task are simply required to make the submission as soon as possible, so that the processing latency is small enough to be neglected. Due to the specific application scenario, two different inference methodologies could be adopted.

Fixed Period The probe publishes sample tasks with the same type and price. After a fixed period T_0, the number of taken tasks as N is observed.

Random Period The probe publishes sample tasks with the same type and price at moment t_0. After N tasks have been taken(or finished), track down the length of the period T_0 starting from t_0.

For both methodologies, under maximum likelihood estimate, the parameter λ_o is given by $\hat{\lambda}_o = \frac{N}{T_0}$. Proof of the correctness of the inference can be found in Appendix Section A. Further advanced sampling-based inference can be found in [31]. The clock rate for the processing phase λ_p is estimated with similar manner. This time, tasks of a specific type are published and the clock rate overall latency is estimated as: $\hat{\lambda} = \frac{N}{T_0}$. Then λ_p is estimated as: $\lambda - \lambda_o$, where λ_o is the estimation of On-hold clock rate.

3.3.2 Linearity Hypothesis

Without loss of generality, within a certain time interval, the price c and the clock rate for the On-hold phase $\lambda_o(c)$ observes relationship with certain linearity. To provide better enhancement of the tuning strategy, we propose a Linearity Conjecture as following, which is the supporting property for strategy in Section 4.2. (The concrete values of the linearity between c and $\lambda_o(c)$ does not affect the design of tuning strategy.)

HYPOTHESIS 1 (LINEARITY). *There exists constant values k and b, such that the rate λ_c and price c follows $\lambda_o(c) = k \cdot c + b$.*

The experiment part gives an empirically justifies this conjecture.

4. TUNING STRATEGIES

In this section, the *H-Tuning* problem is defined in the first place. Then the tuning strategies are developed according to three different scenarios.

4.1 Problem Definition

DEFINITION 2 (LATENCY TARGET). *A* Latency Target L^* *is a stochastic objective function for the tuning problem.*

Specific instantiation of L^* will be presented in each scenario.

DEFINITION 3 (H-TUNING PROBLEM). *Given a set of atomic tasks $T = \{t_1, t_2, \ldots, t_N\}$ with size N, a discrete budget B, find an optimal budget allocation strategy so that* Latency Target L^* *is minimized, without exceeding the budget B.*

4.2 Scenario I - Homogeneity

4.2.1 Scenario Description

Scenario I is the most fundamental case. In this scenario, the system is provided with a set of identical (in terms of difficulty) atomic tasks, which require the same number of running repetitions. All these atomic tasks is published simultaneously, and completed when all the tasks are solved for the required repetitions. A fixed budget is given at the very beginning and the system is to come up with the budget allocation for each atomic task before publishing them to the platform. The budget allocation is made to minimize the expected latency of all the atomic tasks being solved.

4.2.2 Tuning Strategy for Scenario I

The overall latency of all tasks being solved is equivalent to the maximum value of every single task's latency. Specifically, this is defined as $L^* = L(T) = \max\{L(t_i)|i = 1, 2, \ldots, N\}$. As is stated earlier, the latency for each repetition is composed of two phases: the on-hold phase (Phase 1) and the processing phase (Phase 2). The latency of both phases follows an exponential distribution with parameters of λ^o and λ^p. The value of λ^o is determined by the allocated payment with a constant market condition, and the value of λ^p is determined simply by the nature. While our objective is to minimize the overall latency, the budget allocation does not affect the processing latency. Because of the identical nature of the processing time for all the tasks, the minimization of the on-hold latency leads to the minimum latency as well. Therefore, for Scenario I, the objective is changed to the minimization of the expected latency of the on-hold phase. In the remaining part of this section, unless otherwise specified, we use the term "expected latency" referring to the expected latency in Phase 1. Before giving the optimal solution of the budget allocation problem for Scenario I, we introduce the following Lemmas and Theorems.

LEMMA 1. *Given two identical atomic tasks t_1 and t_2, both requiring to be run one round, a fixed budget of B unit payment, allocating both t_1 and t_2 with $\frac{B}{2}$ (or if B is odd, allocating these two atomic tasks with $\lfloor \frac{B}{2} \rfloor$ and $\lfloor \frac{B}{2} \rfloor + 1$) unit payments leads to the minimum expected latency of completing t_1 and t_2.*

PROOF. Please refer to Appendix Section B □

Then, Lemma 2 shows that for one atomic task with multiple repetitions, allocating budget evenly to each repetition will minimize the expected latency.

LEMMA 2. *For atomic task t which needs to be run m repetitions, and a fixed budget of B unit payment, allocating B evenly to each repetition of t leads to the minimum expected latency.*

PROOF. Please refer to Appendix Section C □

With the above two lemmas, we have Theorem 1, which produces the budget allocation plan to minimize the expected latency.

THEOREM 1. *Given two identical atomic tasks which require to be run for the same number of times and a fixed budget of B, allocating the budget evenly to each repetition of all the atomic tasks leads to the minimum expected latency.*

PROOF. Please refer to Appendix Section D. □

directly leads to the optimal budget plan, whose operations are shown in Algorithm 1. As the optimal solution is obtained analytically, EA is conducted with $O(1)$ time complexity.

4.3 Scenario II - Repetition

In this section, we take one more step forward: despite the identical difficulty, the tasks require different running repetitions.

Algorithm 1: Even Allocation (EA)

Input: budget B, atomic task set $T = t_1, t_2, \ldots, t_N$, m required repetition rounds
Output: allocation of payment $P = p_1, p_2, \ldots, p_N$

1 **if** $B \leq m * N$ **then**
2 **return** the budget is not enough;
3 **else**
4 $\delta = \lfloor B/mN \rfloor$ and each repetition of all the atomic tasks is allocated with δ unit payment;
5 $\gamma = \lfloor (B \bmod mN)/N \rfloor$ and select γ repetitions from each atomic task randomly. Increase the payment for the selected rrepetitions by one unit;
6 $\sigma = (B \bmod mN) \bmod N$ and select σ repetitions from σ random atomic tasks whose payment is not increased in the previous step. Increase the payment of the selected repetition rounds by one unit;

4.3.1 Getting the Expected Latency

As the tasks require different number of running repetitions, the closed form of overall latency's pdf will become intractable when tasks come with large quantity. Thus, it's impossible to get the deterministic optimal solution. To address this challenge, the overall latency is processed approximately, based on which the optimal budget plan is derived. Specifically, tasks are grouped according to the running repetitions. Then the overall latency is approximated with the sum of latency of all the task groups.

Group of Single Round Given task group g which is composed of atomic tasks t_1,t_2,...,t_n, requiring to be run for single round. According to the definition, the latency of g, which is denoted by $L(g)$, equals to $\max(L(t_1), L(t_2), \ldots, L(t_n))$. Let $x_1 = \min\{(L(t_1), L(t_2), \ldots, L(t_n))\}$, which means the first completion of all the tasks within the group, and then let $x_2 = \min\{\{(L(t_1), L(t_2), \ldots, L(t_n))\} - x_1\}$ which means the second completion of all the atomic tasks within the group, and the like, $x_n = \min L(g) - \sum_{i=1}^{n-1} x_i$, which means the last completion of all the atomic tasks. It can be derived that $L(g) = \sum_{i=1}^{i=n} x_i$. As $x_i \sim exp(\lambda * i)$, $L(g)$ can be regarded as the sum of n exponential variables. Therefore, $L(g) = \sum_{i=1}^{n} \frac{1}{\lambda * i}$.

Group Multiple Rounds Before we turn to the study of the probabilistic model of the task group of multiple repetition rounds, the following lemma is needed to show the probabilistic property of the task which requires multiple running repetitions.

LEMMA 3. *Let t denote an atomic task which needs to be run for k repetition rounds, the latency of t follows Erlang distribution of parameter k and λ, which is $L(t) \sim Erl\{k, \lambda\}$*

PROOF. Please refer to Appendix Section E. □

Now we can get the expected latency of the task group through the following deduction. Suppose we are given a task group g, which is composed of a set of tasks $\{t_1, t_2, \ldots, t_n\}$ and each task is needed to be run for k repetition rounds. Let $L\{g\}$ denote the latency of the task group g, then we can have the following relationship: $L\{g\} = L\{max\{x_1, x_2, \ldots, x_n\}\}$. Let $F(t)$ denote the cumulative distribution function (cdf) and $f(t)$ denote the probability density function (pdf) of the latency of the atomic tasks respectively. Let $F_g(t)$ denote the cumulative distribution function (cdf) and $f_g(t)$ denote the probability density function (pdf) of the latency of the task group respectively. The following relationship can be derived:

$$F_g(t) = F^n(t), \quad f_g(t) = n * F^{n-1}(t) * f(t)$$

With the above relationship, we get get the expression of the ex-

pected latency of the task group as:

$$E\{L(g)\} = \int_0^\infty f_g(t) * tdt = \int_0^\infty n * F^{n-1}(t) * f(t) * tdt$$

According to the conclusion of $lemma$ 3, the latency of the tasks within the task group follow Erlang distribution $Erl(k, \lambda)$, thus the expected latency of the task group is derived as:

$$E\{L(g)\} = \int_0^\infty n * F_E^{n-1}(k, \lambda, t) * f_E(k, \lambda, t) * tdt$$

among which $F_E(k, \lambda, t)$ and $f_E(k, \lambda, t)$ denote the cdf and pdf of the Erlang distribution $Erl(k, \lambda)$.

Approximate Expected Latency As is stated in the previously, the close form of expected latency is intractable when the number of tasks is huge. Thus, we use the sum of the expected latency of all the task groups to approximate the real function. There are two reasons for such approximation: one is that the sum of the expected latency of all the task group lays the upper bound of the expected latency of all the atomic tasks; the other is that the expected latency of all the atomic tasks will decrease while the sum of the expected latency of the task group is going down.

4.3.2 Tuning Strategy for Scenario II

Let $E_{g_1}, E_{g_2}, \ldots, E_{g_n}$ denote the expected latency of the task group $g_1, g_2, \ldots, g_n$. Let $b_{g_1}, b_{g_2}, \ldots, b_{g_n}$ denote the allocated payment of task group $g_1, g_2, \ldots, g_n$. The optimizing problem is defined as $\min \sum_{i=1}^n E_{g_i}\ s.t. \sum_{i=1}^n b_{g_i} \leqslant B$.

A dynamic algorithm is designed as follows to solve such minimization problem. The outer loop of the algorithm increases the task payment from 1 to B' ($B' = B - \sum_{i=1}^n u_i$). Within each loop, it takes $O(n)$ operations to find the optimal payment given the current budget. Apparently, the overall time complexity for algorithm 2 turns out to be $O(nB')$

Algorithm 2: Repetition Algorithm (RA)

Input: budget B, task group $G = g_1, g_2, \ldots, g_n$
Output: allocation of payment $P = p_1, p_2, \ldots, p_n$

```
for i = 1 to n do
    p_i(0) = 1
B' = B - Σ_{i=1}^n u_i;
E_0(0) = Σ_{i=1}^n E_i(P_i(0));
for x = 1 to B' do
    E_0(x) = min{E_0(x-1),
    min{E_0(x-u_i) - [E_i(p_i) - E_i(p_i+1)] | u_i ≤ x}};
    if
    E_0(x-1) ≤ {E_0(x-u_i) - [E_i(p_i) - E_i(p_i+1)] | u_i ≤ x}
    then
        θ =
        argmin_i {E_0(x-u_i) - [E_i(p_i) - E_i(p_i+1)] | u_i ≤ x};
        p_θ(x) = p_θ(x-1) + 1;
∀i = 1, ..., n, p_i = p_i(b)
```

4.4 Scenario III - Heterogeneous

In Scenario III, the tasks are heterogeneous (in terms of difficulty) and need to be run for different numbers of repetitions. While dealing with the latency of first two scenarios, we only take the latency of Phase 1 into account, whose reasons are two fold: the first one is that the payment does not change the latency of Phase 2, the second one is that the latency of Phase 2 is identical for all the atomic tasks since all the tasks are homogeneous in terms of task nature. However, these properties no longer hold in Scenario III as different tasks require different processing time.

For this scenario, some tasks are easier to solve, which produce smaller processing latency, while others are harder to solve, which lead to longer processing latency. As a result of such character, the previous tunning strategies do not apply well to the current problem, as the tuning result may be jeopardized by the tasks whose processing latency is significantly larger than others'. One extreme situation is that the latency of Phase 2 of some atomic tasks is so long that the overall latency of completing all the atomic tasks will be approximately equal to the expected latency such atomic task. We call such kind of atomic tasks as "most difficult task". It is obvious that such type of atomic tasks generate stronger influence to the overall latency than the others.

In order to relieve the delaying effect caused by the "most difficult tasks", we make the following adaption to the tuning strategy. For previous tuning strategies, only the latency of Phase 1 is considered. While in Scenario III, two objectives will be minimized simultaneously: one objective is still the latency Phase 1, the other one is the latency of the "most difficult task", which is equivalent to the largest expected latency of all the atomic tasks. The reason of introducing the first objective is the same as previous scenarios, which is the allocation of payment only changes the latency of Phase 1, while the reason of introducing the second objective is to confine the delay effect caused by the "most difficult task". Here the second objective serves as the penalty function to avoid the appearance of the situation where the latency of some atomic tasks is significantly longer than that of others'. One more point needs to be clarified is that we can't simply minimize the second objective because the minimization of the latency of the "most difficult work" doesn't necessarily lead to minimum latency of completing all the atomic tasks.

Formally, the objective function is defined as follow. Let $G = \{g_1, g_2, \ldots, g_n\}$ denote the task group (the grouping operation is performed to all the atomic tasks so that the tasks of identical type and repetition fall into the same group, which is slightly different from Scenario II). Let $L^1(g_i)$ and $L^2(g_i)$ denote the latency of Phase 1 and Phase 2 of g_i respectively. Objective 1 is the expected latency of Phase 1 of all the atomic tasks, which is denoted by O_1 and $O_1 = E\{L^1(G)\}$. Objective 2 is the sum of the expected latency of Phase 1 and Phase 2 of the most difficult atomic tasks, which is denoted by O_2 and $O_2 = max\{E\{L^1(g_i)\} + E\{L^2(g_i)\} | i = 1, \ldots, n\}$. Given the budget of B unit payment and let p_i denote the payment allocated to group g_i, the optimizing problem is defined as: $\min\{O_1, O_2\}\ s.t.$ $\sum_{i=1}^n p_i \leq B$. Here, we adopt a "Compromise strategy" to solve the above two objective optimization problems. Firstly, the "Utopia Point" (UP) is calculated, which refers to the point where both objectives are optimized independently under the given constraints. In the second place, the "Closeness" (CL) is defined as the first order distance between the objective point OP and UP. The "Closeness" is minimized under the given constrains, and the corresponding solution will serve as the optimal solution. The definition of "UP", "OP", and "CL" are formally presented as follows.

DEFINITION 4 (UTOPIA POINT). *Let $O_1^* = \{\min O_1 : s.t. \sum_{i=1}^n p_i \leq B\}$ and $O_2^* = \{\min O_2 : s.t. \sum_{i=1}^n p_i \leq B\}$. The Utopia Point is defined as $UP = (O_1^*, O_2^*)$.*

DEFINITION 5 (OBJECTIVE POINT). *Let O_1 and O_2 denote the objective value of the current payment allocated to each task group. The Objective Point is the two dimensional position determined by {O_1,O_2}.*

DEFINITION 6 (CLOSENESS). *The Closeness equals to the first order distance between UP and OP: $CL = \|OP - UP\|$.*

Here, the optimal budget plan is equivalent to the minimization of the following problem: $\min CL : s.t. \sum_{i=1}^{n} p_i \leq B$. Such a problem can be optimally solved with dynamic programming, whose procedures are shown with Algorithm 3. Similar with algorithm 2, the dynamic programming runs with $O(nB')$ iterations to achieve the optimal solution.

Algorithm 3: Heterogeneous Algorithm (HA)

Input: budget B, task group $G = g_1, g_2, \ldots, g_n$
Output: allocation of payment $P = p_1, p_2, \ldots, p_n$

for $i = 1$ *to* n **do**
  $p_i = 1$
$B' = B - \sum_{i=1}^{n} u_i$;
$CL_0 = \|OP_0 - UP\|$;
for $x = 1$ *to* B' **do**
  $CL_x = \min\{CL_{x-1}, \min\{CL_{x-u_i}(+ + p_i(x - u_i))$
  $|u_i \leq x, i = 1, \ldots, n\}\}$;
  if CL_{x-1}
  $\leq \{CL_{x-u_i}(+ + p_i(x - u_i))|u_i \leq x, i = 1, \ldots, n\}$ **then**
    $\theta = \underset{i}{\operatorname{argmin}}\{CL_{x-u_i}(+ + p_i(x - u_i))|u_i \leq x, i = 1, \ldots, n\}$;
    $p_\theta(x) = p_\theta(x - 1) + 1$;
$\forall i = 1, \ldots, n, p_i = p_i(b)$

5. EXPERIMENTS

We extensively evaluated our model and optimization techniques, and report on the results here. While the gold standard is performance on a real platform, we can exercise greater control and thereby get a better empirical understanding of our system through simulation. Therefore, we did both. We report first on simulation results with synthetic data, and then on jobs executed on Amazon Mechanical Turk.

5.1 HPU Traits Testing with Synthetic Data

5.1.1 Experiments Settings

We conduct six sets of experiment for each Scenario. The first four sets are linear model based, which aim to verify the effectiveness of the tuning strategy under the linear Hypothesis, and the last two sets are nonlinear model based, which aim to test the robustness of the tuning strategy. For the linear model based experiment, the model parameters are set as $\lambda^o = p + 1$, $\lambda^o = 10p + 1$, $\lambda^o = 0.1p + 10$, $\lambda^o = 3p + 3$. For the nonlinear part, the parameters are set as $\lambda^o = 1 + p^2$ and $\lambda^o = \log(1 + p)$. The total number of task is set to be 100 uniformly for each set of the experiment and the budget varies from 1000 to 5000.

Homogeneity All tasks call fro 5 repetitions. As the difficulty of the tasks are identical, the clock rate λ_p for the processing latency is uniformly set to be 2.0. Since the optimistic solution is produced by the *even allocation* (*algorithm 1*), biased allocation strategies are adopted as the baseline comparison. Instead of allocating the budget evenly, the biased method gives more payment to one half of the tasks, while less payment to the other. Specifically, half of the tasks are randomly selected as "the prior group" which take up α ($\frac{1}{2} < \alpha < 1$) of the total budget ($\alpha = \frac{1}{2}$ leads to the even allocation), and the remaining tasks get the $1-\alpha$ of the total budget. The value of the alpha is set to be 0.67, and 0.75 in our experiment.

Repetition The tasks are divided equally into two groups: one group is of 3 repetitions for each task, while the other group is of 5 repetitions for each of the tasks. Still, λ_p is uniformly set to be 2.0 due the identical setting of the difficulty. Two baseline methods are chosen as for the comparison. The first method is called *task-even* allocation, which gives identical price to each task, then every task allocate the total budget evenly to each of its repetitions. Therefore, repetition price for group 2 is 60% of that of group 1. The second one is called *rep-even* allocation, which gives identical price to each repetition of all the tasks. So the total price for the tasks in the group 1 is 60% of that of group 2.

Heterogeneous The tasks are dived into two groups: task in the first group call fro three repetitions, while tasks in the second group call fro five repetitions. Then λ_p is set to be 2.0 and 3.0 these two groups, respectively. Same with scenario ii, the *rep-even* and *task-even* are chosen the baseline methods for the comparison.

We also conduct experiments with different settings of the budget, task amount, repetitions, and difficulty. However, there's no significant variance between different settings. Therefore, we simply demonstrate the results of the above setting for further analysis.

5.1.2 Results Summary

From the experiment results, the optimal solution outperforms the comparisons in terms of latency in every cases. For results of scenario 1 (homo), the "*bias_1*" produces slightly better performance than "*bias_2*". This is because *bias_2* is more biased (the value of α is larger) than *bias_1*. Such phenomenon further verified our conclusion that even allocation leads to the optimal budget plan for the Scenario 1. Besides, we can find that although the optimal solution of the Scenario 1 is designed based on the linear hypothesis, it still works for the nonlinear cases (*homo(e)* and *homo(f)*), which can be partially explained by the varying range of the payment: the task price varies form 1 to 9. For such relatively low prices, the non-linear relationship can be linearly approximated quite well. We can further find that the optimal results are relatively close to the comparison in case (b) and (c), for all the scenarios. For case (b), such phenomenon can be caused the large value of the linear coefficient ($lambda = 10p + 1$). When the linear coefficient is large, the on-hold clock rate is sensitive to the change of price. When price grows, the clock rate increases much more faster, making the oh-hold latency decrease to a low level with a relatively lower price. In this situation, the overall latency will be mostly determined by the processing phase. Similar phenomenons can be observed for case (e) in each of the scenarios, where overall latency reduces sharply for the initial prices, and soon get to a stable level. While case (c) is another extreme, where λ is fairly insensitive to the price changes. In this situation, and the latency is largely determined by the initial setting of on-hold and processing phase, and price does little to change it.

Finally, we can summarize the findings of the synthetic experiment as follows: 1) the optimal tuning strategy is robust to nonlinearity. The unit price for each task is usually small, therefore the linearity hypothesis holds for normal cases. 2) The optimal tuning strategy is sensitive to the *price-lambda* relationship: when $lambda$ is sensitive to the change of price, the on-hold latency drops sharply with the growing price. Then the overall latency is determined by the processing time and it's unnecessary to keep on increasing the price.

5.2 Tuning Tasks on Amazon MTurk

5.2.1 Experiments Settings

We create a set of image filtering tasks as the atomic tasks: we first present the workers an image with the exact number of the dots on it, then a set of images are presented to the "workers" and they are required to estimate the number of dots on each image. Based on the estimation, "workers" are expected to filter out the ones who have dots less than a given threshold. Under such settings, the cognitive abilities of "recognizing" and "counting" are utilized, and the

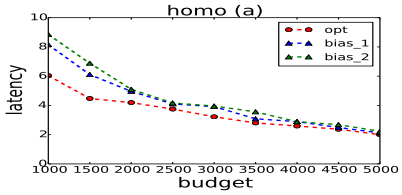

(a) Homogeneous($\lambda = 1 + p$)

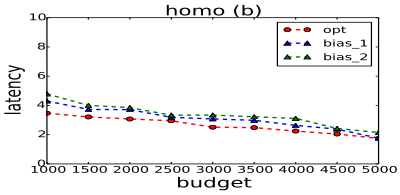

(b) Homogeneous($\lambda = 10p + 1$)

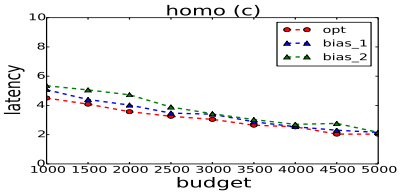

(c) Homogeneous($\lambda = 0.1p + 10$)

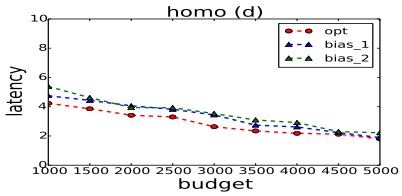

(d) Homogeneous($\lambda = 3p + 3$)

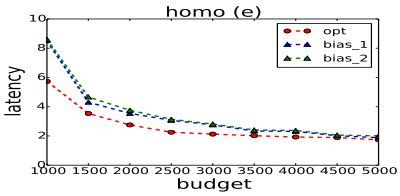

(e) Homogeneous($\lambda = 1 + p^2$)

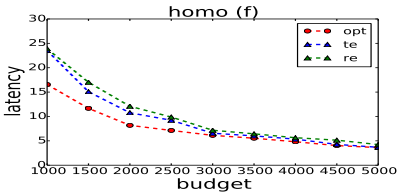

(f) Homogeneous($\lambda = \log 1 + p$)

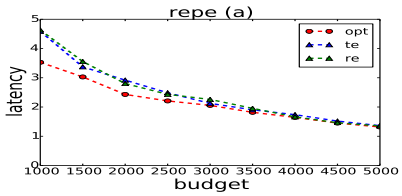

(g) Repetition($\lambda = 1 + p$)

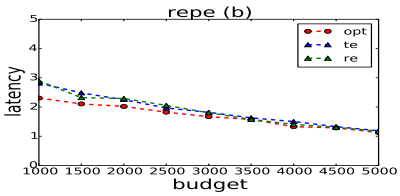

(h) Repetition($\lambda = 10p + 1$)

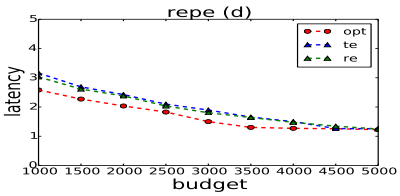

(i) Repetition($\lambda = 0.1p + 10$)

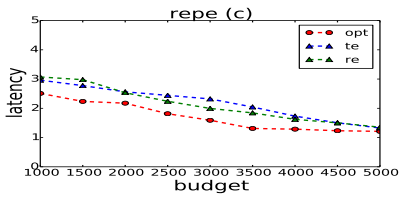

(j) Repetition($\lambda = 3p + 3$)

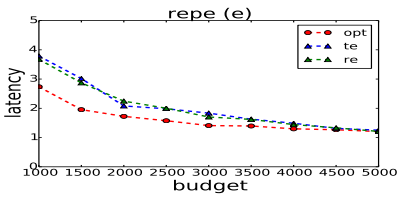

(k) Repetition($\lambda = 1 + p^2$)

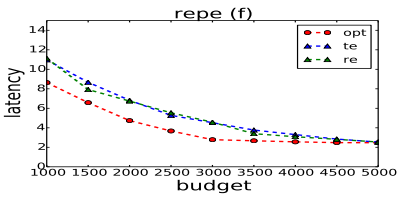

(l) Repetition($\lambda = \log 1 + p$)

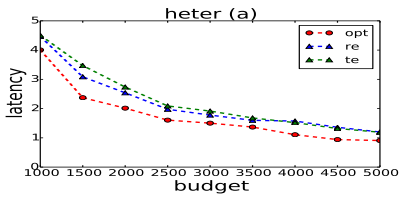

(m) Heterogeneous($\lambda = 1 + p$)

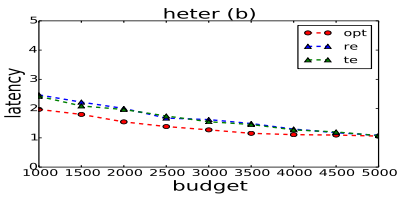

(n) Heterogeneous($\lambda = 10p + 1$)

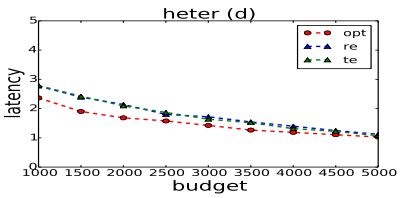

(o) Heterogeneous($\lambda = 0.1p + 10$)

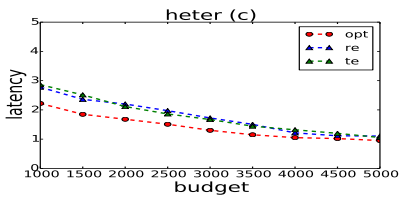

(p) Heterogeneous($\lambda = 3p + 3$)

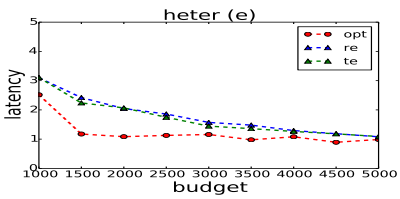

(q) Heterogeneous($\lambda = 1 + p^2$)

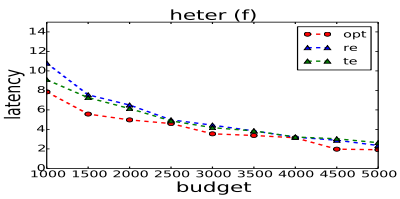

(r) Heterogeneous($\lambda = \log 1 + p$)

Figure 2: Experiments on Synthetic Data

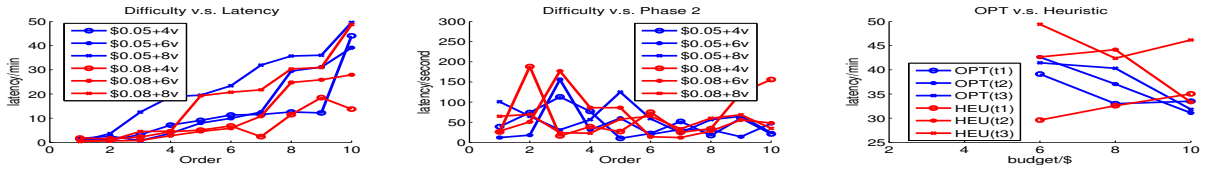

(a) Difficulty v.s. Phase 1 (b) Difficulty v.s. Phase 2 (c) OPT v.s. Heuristic

Figure 5: Experiments Results II

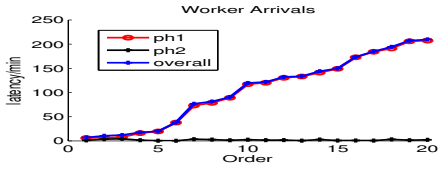

Figure 3: Worker Arrival Moments

Figure 4: Money v.s. Latency

task is finished by presenting a set of binary voting(clicking on the checkbox). In addition, the "workers" receive their rewards when the provided answers are correct. We control the difficulty or type of tasks by varying the images given in a single tasks.

Our work focuses on tuning upon budget allocation and real time latency, thus we purposely design the experiment simple enough and avoid setting any worker qualifications and inter-rater agreement. In fact, in real scenarios, the atomic tasks on lowest level are just the same as the experimental tasks: comparing items, screening out candidates and simple ranking.

5.2.2 *Results Summary*

Firstly, in Fig 3 we present the general behavior of the worker appearance and the latency of processing time. We issue image filtering tasks with 1 unit reward($0.05), and collect the first 20 arrivals. As shown in Fig 3, the arrival epochs of the workers exhibit linearity, indicating the suitability of the Poisson Process Model, while the latency of the second phase fluctuates in a small range.

Then we examine the effect of varying the rewards: we vary the reward on a single task from $0.05 to $0.12, while for each task we require 10 repetitions. The results can be found in Fig 4, where it is obviously that the increase on rewards incurs shorter latencies. According to the methodologies introduced in Section 3.3, we obtain the corresponding parameters(s^{-1}), $\lambda_1 = 0.0038, \lambda_2 = 0.0062, \lambda_3 = 0.0121, \lambda_4 = 0.0131$, which supports the Linearity Hypothesis proposed in Section 3.3.2.

In the sequel, we present the results of examining the effect of varying the type of the tasks: we vary the internal binary voting number from 4 to 8. Such change of difficulty results in the decrease of the coming rate (see Fig 5(a)), and the increase of the average processing time, which is shown in Fig 5(b). We then evaluate our proposed algorithms on Amazon MTurk, especially under Scenario II and III. Namely, 3 types of tasks are published with different repetition requirement: 10 for $t1$, 15 for $t2$ and 20 for $t3$. The total budgets are also varied from $6 to $10. We compare our algorithms(OPT) with the heuristic where each type receives same payment under both two scenarios. Results can be found in Fig 5(c), where the lower latency of OPT shows the effectiveness of our algorithms. Note that at each budget, the OPT successfully avoids yielding the longest latency among the three tasks.

6. CONCLUSION

In this paper, we address the problem of tuning the modularized human computation, so that the latency in real clock time could be minimized. The difficulty of such problem arises in the stochastic behavior of the latency of the HPU. To address this challenge, we theoretically and practically propose that appearance of the crowd "workers" follows a Poisson Process, whose parameter differs at different budget levels and types of atomic tasks. Then we formally propose the *H-Tuning Problem* to optimize the expected latency of the longest task. Moreover, under three most general scenarios on crowd-powered applications, advanced strategies are designed to cope with the *H-Tuning Problem*. Finally, a series of experiments conducted on both simulated data and real commercial platform observe the effectiveness of the proposed model and strategies. To conclude, the crowdsourced human computation is now equipped with primitive tuning ability in terms of running time.

7. REFERENCES

[1] J. Davis, J. Arderiu, H. Lin, Z. Nevins, S. Schuon, O. Gallo, and M. Yang, "The hpu," in *CVPRW 2010*.

[2] *Human Computation*.

[3] D. A. Grier, "The math tables project of the work projects administration: The reluctant start of the computing era," *IEEE Ann. Hist. Comput.*

[4] A. Parameswaran and N. Polyzotis, "Answering queries using humans, algorithms and databases," in *CIDR 2011*.

[5] M. J. Franklin, D. Kossmann, T. Kraska, S. Ramesh, and R. Xin, "Crowddb: answering queries with crowdsourcing," in *SIGMOD 2011*.

[6] A. Marcus, E. Wu, D. Karger, S. Madden, and R. Miller, "Human-powered sorts and joins," *VLDB 2011*.

[7] A. G. Parameswaran, H. Garcia-Molina, H. Park, N. Polyzotis, A. Ramesh, and J. Widom, "Crowdscreen: algorithms for filtering data with humans," in *SIGMOD 2012*.

[8] P. Venetis, H. Garcia-Molina, K. Huang, and N. Polyzotis, "Max algorithms in crowdsourcing environments."

[9] S. Guo, A. G. Parameswaran, and H. Garcia-Molina, "So who won?: dynamic max discovery with the crowd," in *SIGMOD Conference 2012*.

[10] S. B. Davidson, S. Khanna, T. Milo, and S. Roy, "Using the crowd for top-k and group-by queries," in *ICDT 2013*.

[11] X. S. Yang, R. Cheng, L. Mo, B. Kao, and D. W. Cheung, "On incentive-based tagging," in *ICDE 2013*.

[12] J. C. Zhang, L. Chen, H. V. Jagadish, and C. C. CAO, "Reducing uncertainty of schema matching via crowdsourcing," *VLDB 2013*.

[13] J. Wang, G. Li, T. Kraska, M. J. Franklin, and J. Feng, "Leveraging transitive relations for crowdsourced joins," in *SIGMOD 2013*.

[14] I. Lotosh, T. Milo, and S. Novgorodov, "Crowdplanr: Planning made easy with crowd," in *ICDE 2013*.

[15] Y. Amsterdamer, Y. Grossman, T. Milo, and P. Senellart, "Crowd mining," in *SIGMOD 2013*.

[16] J. Wang, S. Faridani, and P. Ipeirotis, "Estimating the completion time of crowdsourced tasks using survival analysis models," *CSDM 2011*.

[17] S. Faridani, B. Hartmann, and P. G. Ipeirotis, "What's the right price? pricing tasks for finishing on time." in *Hcomp 2011*.

[18] W. Mason and D. J. Watts, "Financial incentives and the "performance of crowds"," ser. HCOMP 2009.

[19] T. Yan, V. Kumar, and D. Ganesan, "Crowdsearch: exploiting crowds for accurate real-time image search on mobile phones," in *MobiSys 2010*.

[20] X. Liu, M. Lu, B. C. Ooi, Y. Shen, S. Wu, and M. Zhang, "Cdas: a crowdsourcing data analytics system," *VLDB 2012*.

[21] C. C. CAO, J. She, Y. Tong, and L. Chen, "Whom to ask? jury selection for decision making tasks on micro-blog services," *VLDB 2012*.

[22] P. G. Ipeirotis, F. Provost, and J. Wang, "Quality management on amazon mechanical turk," in *Proceedings of the ACM SIGKDD Workshop on Human Computation*.

[23] R. Boim, O. Greenshpan, T. Milo, S. Novgorodov, N. Polyzotis, and W. C. Tan, "Asking the right questions in crowd data sourcing," in *ICDE 2012*.

[24] J. P. Bigham, C. Jayant, H. Ji, G. Little, A. Miller, R. C. Miller, R. Miller, A. Tatarowicz, B. White, S. White, and T. Yeh, "Vizwiz: nearly real-time answers to visual questions," in *UIST 2010*.

[25] G. Pickard, W. Pan, I. Rahwan, M. Cebrian, R. Crane, A. Madan, and A. Pentland, "Time-critical social mobilization," *Science*, 2011.

[26] M. S. Bernstein, J. Brandt, R. C. Miller, and D. R. Karger, "Crowds in two seconds: Enabling realtime crowd-powered interfaces," in *UIST 2011*.

[27] M. S. Bernstein, D. R. Karger, R. C. Miller, and J. Brandt, "Analytic methods for optimizing realtime crowdsourcing," *arXiv 2012*.

[28] P. Minder, S. Seuken, A. Bernstein, and M. Zollinger, "Crowdmanager - combinatorial allocation and pricing of crowdsourcing tasks with time constraints," in *ACM-EC 2012*.

[29] Y. Gao and A. Parameswaran, "Finish them!: Pricing algorithms for human computation," in *VLDB 2014*.

[30] J. Ross, L. Irani, M. Silberman, A. Zaldivar, and B. Tomlinson, "Who are the crowdworkers?: shifting demographics in mechanical turk," in *CHI 2010*.

[31] I. Basawa and B. Rao, *Statistical inference for stochastic processes*, ser. Probability and mathematical statistics. Academic Press, 1980.

APPENDIX

A. INFERENCE OF PARAMETER λ

Fixed Period The likelihood function of parameter λ is $L = \lambda_N \exp[-\lambda \sum_{k=1}^{N} T_k] \exp[-\lambda(T_0 - t_n)] = \lambda_N \exp[-\lambda T_0]$. To maximize the likelihood, the ML estimation of λ is derived as $\hat{\lambda} = N/T_0$, which is unbiased according to Rao-Blackwell Theorem.

Random Period Suppose each worker appears at the epochs $0 < t_1 < t_2 \ldots < t_N$, we could obtain the likelihood function: $L = \lambda^N e^{-\lambda T_0}$. Thus to maximize the likelihood, the ML estimation of λ is given by $\hat{\lambda} = \frac{N}{T_0}$(same as the one in Fixed Period). To remove the bias, further the parameter can be updated: $\tilde{\lambda} = ((N-1)N)\hat{\lambda}$.

B. PROOF OF LEMMA 1

PROOF. As illustrated in the previous section, Phase 1 of both t_1 and t_2 follow exponential distribution, whose parameters are denoted as $\lambda^o_{t_1}$ and $\lambda^o_{t_2}$, and according to Hypothesis I in Section 3.3, when allocating t_1 with x unit payment and t_2 with $B-x$ unit payment, $\lambda^o_{t_1} = kx$ and $\lambda^o_{t_2} = k(B-x)$. (k is the constant coefficient) With the two parameters defined above, we can derive the following relationship: $E(\{t_1, t_2\}) = E\{\max\{L^o(t_1), L^o(t_2)\}\} = \frac{\lambda^o_{t_1}+\lambda^o_{t_2}}{\lambda^o_{t_1}\lambda^o_{t_2}} - \frac{1}{\lambda^o_{t_1}+\lambda^o_{t_2}}$ As $\lambda^o_{t_1} = kx$ and $\lambda^o_{t_2} = k(B-x)$, $E_(t_1, t_2)$ turns out to be a convex function and reaches its minimum point when $x = \frac{B}{2}$(or $\lfloor B/2 \rfloor$ if B is odd). Hence, allocating t_1 and t_2 with $\frac{B}{2}$ unit payments leads to the minimum expected latency. □

C. PROOF OF LEMMA 2

PROOF. Let $\{p_1, \ldots, p_m\}$ denote the payment allocated to each repetition of atomic task t, and $\{\lambda^o_1, \ldots, \lambda^o_m\}$ denote the exponential parameter of each repetition. It is obvious to see that $\sum_{i=1}^{k} p_i = B$, and based on Hypothesis 1, $\lambda^o_i = k * p_i$. With above information, we can derive the expected latency of t as follows.
$E\{L(t)\} = \sum_{i=1}^{m} 1/\lambda^o_i = \sum_{i=1}^{m} 1/kp_i$. Since $E\{L(t)\} = \sum_{i=1}^{m} 1/kp_i \leq B^2/km$, and the equality is established iff. $\forall i \in \{1, \ldots, m\}$,$x_i = B/m$. Therefore, we come to the conclusion that allocating the budget evenly to each repetition of the atomic task leads to the minimum expected latency. □

D. PROOF OF THEOREM 4.2.2

PROOF. This theorem is proved with mathematical induction. Firstly, the theorem holds when both atomic tasks are run exactly once, which is a direct result of Lemma 1. Then, we prove that the theorem still holds when the repetitions increases to $n+1$ on condition that it hold with repetition equals to n. Suppose we have a two identical tasks t_1 and t_2, which require n reps. Now we allocate each rep with x unit payment.Based on our presumption, this budget allocation leads to the minimum expected latency. Let H denotes the completion of task t_1, and the completion of both tasks as $\max\{H, H\}$. Now we increase the repetitions of both tasks to $n+1$, and the budget to $2(n+1)x$. Suppose we have a better budget allocation which outperforms allocating the budget evenly. It is trivial to see that one task will be allocated with more payment and the other atomic task will be allocated with less payment. At the same time the payment for each repetition of the same atomic task remains identical. Let I denote the completion of the first n repetitions of t_1 and i denote the event of the completion of the last repetition of t_1. So the completion of t_1 is denoted as $\{I+i\}$. Let J denote the completion of the first n repetitions of t_2 and j denote the completion of the last repetition of t_2. The completion of t_2 can be denoted as $J+j$. Similarly, when allocating the budget evenly to each repetition of both tasks, we use H to denote the first n repetitions of the atomic task and h to denote the last repetition of the atomic task, and the completion of t_1 (or t_2) is denoted as $\{H+h\}$. Then, the completion of both tasks with the assumed optimal budget allocation is denoted by $\max\{\{I+i\}, \{J+j\}\}$ and the completion of both tasks with the evenly allocated budget is denoted as $\max\{\{H+h\}, \{H+h\}\}$. Here, we will have the flowing relationship: $E\{\max\{\{I+i\}, \{J+j\}\} = (E\{I+i\} + E\{J+j\}) - E\{\min\{\{I+i\}, \{J+j\}\}\}$, and
$E\{\max\{\{H+h\}, \{H+h\}\}\} = (E\{H+h\} + E\{H+h\}) - E\{\min\{\{H+h\}, \{H+h\}\}\}$. As $E\{I+i\} + E\{J+j\} = n/(k(x-\varepsilon)) + n/(k(x+\varepsilon)) \leq 2n/kx = E\{H+h\} + E\{H+h\}$, we can get the result that $E\{\max\{\{I+i\}, \{J+j\}\}\} \geq E\{\max\{\{H+h\}, \{H+h\}\}\}$. This shows that the theorem still holds when the repetitions increases to $n+1$, which prove the theorem to be truth. □

E. PROOF OF LEMMA 3

PROOF. Let r_i denote the ith repetition of task x. According to Lemma 1,$\forall r_i \in \{r_1, \ldots, r_n\}$, $L(r_i)$ follows exponential distribution of the same parameter λ. So $L(x) = L(\sum_{i=1}^{k} r_i)$. This meets the requirement of Erlang distribution and makes $L(x) \sim Erl(k, \lambda)$. □